arXiv:1611.06633v1 [math.RA] 21 Nov 2016

Two new direct linear solvers in the QR family

Michael F. Zimmer

November 22, 2016

Abstract

Two new algorithms for computing a direct solution to a system of linear equations are presented. A variation on the orthonormalization step in the usual QR method allows one to bypass subsequent elimination or matrix inversion steps. The algorithms are formulated with respect to row and column spaces, and produce a generalized inverse in each case. They may also be formulated in an online manner, so that the solution may be computed as the data for the equations become available.

AMS classification: 15A06, 15A09

Key Words: system of linear equations; row space; column space; generalized inverse; minimum norm; least squares; direct solution; online solution

1 Introduction

The problem of solving a system of linear equations is widespread across mathematics, science and engineering. Taking there to be m equations in n unknowns, the equations will be written as $Ax = b$, where $A = (a_{ij})$ is the coefficient matrix, $x = (x_j)$ is the vector of unknowns, and $b = (b_i)$. The indices are $i = 1, \ldots, m$ and $j = 1, \ldots, n$, and the entries in A, x and b belong to a field $\mathcal{F}$ which is taken to be $\mathcal{C}$.

When the nullity of A is non-zero, the inverse to A does not exist, and it will no longer possible to write the solution in the cogent form $x = A^{-1}b$.

Nevertheless, in this case one may still use the generalized inverse (G) to express the particular and homogeneous parts of the solution as:

$$\begin{aligned} x_p &= Gb \\ x_h &= Py \\ P &= 1_n - GA \end{aligned}$$

where $Ax_p = b$, $Ax_h = 0$, 1_n is an n-by-n identity matrix, P is a null space projection operator, and $y \in \mathcal{F}^n$. By definition, $Py \in \mathcal{N}[A]$, where $\mathcal{N}[A]$ is the null space of A.

The methods to be discussed will be seen to bear a similarity to the family of QR factorization methods for solving $Ax = b$ (i.e., the QR, RQ, LQ and QL methods)[12]. Those methods involve, for example, making the substitution $A = QR$ in $Ax = b$, and then solving for x. In contrast, the methods proposed herein involve a factorization of A while it is in the context of the equation $Ax = b$. They'll be referred to as "in situ" methods, to emphasize this point. In contrast, since the QR methods involve a factorization of A done externally to $Ax = b$, they might instead be called "ex situ" methods. Finally, it will be shown that in situ approaches allow for an "online" solution, which is not prima facie possible with the QR methods.

2 Definitions

The two main algorithms to be presented will involve orthonormalizing the row (column) vectors of the coefficient matrix A, by left (right) multiplying A with a series of nonsingular matrices M_s ($s = 1, 2, \dots$). It's not necessary for the M_s to implement elementary row (column) operations, but it's convenient to view them that way. In the course of applying an orthonormalization procedure (e.g., Gram-Schmidt [7, 12, 16, 15]) to the rows or columns of A, it may be that a zero vector will appear, since there may be linear dependencies amongst those vectors. If they do appear, they will be left in place in A and thereafter ignored. The result will be that the vectors of the transformed A will form a quasi-orthonormal list of vectors, as defined next.

Definition: A **quasi-orthonormal** list of vectors consists of vectors with (Euclidean) norm equal to 1 or 0. Those vectors of norm 1 are mutually orthogonal.

It will later be convenient to use a matrix which is essentially a p-by-p identity matrix, but with some of the 1s replaced by 0s. The row (or column) indices of the 1s will comprise the set S. Note that $S \subseteq \{1, 2, \ldots, p\}$. Thus an **index matrix** is defined as:

$$(\mathcal{I}_p^S)_{ij} = \begin{cases} 1 & \text{if } i = j \text{ and } i \in S, \\ 0 & \text{otherwise.} \end{cases}$$

where i, j each range from 1 to p. Note that if $|S| = p$, then $\mathcal{I}_p^S$ is just 1_p.

3 Row-Based In Situ Approach

The first step is to attempt to orthonormalize the rows of A, using the approach outlined in the previous section. The result is that $Ax = b$ becomes

$$(\cdots M_2 M_1)Ax = (\cdots M_2 M_1)b$$
$$A'x = b'$$

where $M = (\cdots M_2 M_1)$, $A' = MA$, and $b' = Mb$ [14]. The row vectors are now a quasi-orthonormal list of vectors.

Next define the set S to contain the indices of the rows of A' that are non-zero. It is then easy to verify

$$A' \ (A')^* = \mathcal{I}_m^S \tag{1}$$
$$\mathcal{I}_m^S \ A' = A' \tag{2}$$
$$\mathcal{I}_m^S \ b' = b' \tag{3}$$

where the superscript * represents a conjugate transpose. A solution will now be claimed and verified. Afterwards, a constructive approach will be taken.

Theorem: A solution to $Ax = b$ is

$$x = (A')^* b' + x_h \tag{4}$$

where the variable definitions from this section are used.

Proof. The proposed solution is verified by substitution into Ax

$$\begin{aligned} Ax &= A(A')^* b' + Ax_h \\ &= M^{-1}A'(A')^* b' \\ &= M^{-1}\mathcal{I}_m^S b' \\ &= M^{-1}b' \\ &= b \end{aligned}$$

□

Having verified an unmotivated solution to $Ax = b$, a solution will now be derived by constructive means.

Lemma: An arbitrary vector $x \in \mathcal{F}^n$ may be expressed as

$$x = (A')^* w + x_h \tag{5}$$

where A is an m-by-n matrix, $w \in \mathcal{F}^m$ and x_h satisfies $Ax_h = 0$.

Proof. Because of the linear transformation $A' : \mathcal{F}^n \to \mathcal{F}^m$, it is known (Thm. 18.3 [3]) that

$$\mathcal{F}^n = \mathcal{R}[(A')^*] \oplus \mathcal{N}[A']$$

where $\mathcal{R}[A]$ is the range of A. This direct sum decomposition may be used to rewrite an $x \in \mathcal{F}^n$ as

$$x = w_1 + w_2$$

where

$$\begin{aligned} w_1 &\in \mathcal{R}[(A')^*] \\ w_2 &\in \mathcal{N}[A'] \end{aligned}$$

However, every vector that is in the range of some matrix Q may be expressed as Qv, for an appropriate v. Thus, for some $w \in \mathcal{F}^m$, w_1 may be written as

$$w_1 = (A')^* w$$

Also, because M is nonsingular, $\mathcal{N}[A'] = \mathcal{N}[A]$, so that $w_2 \in \mathcal{N}[A]$. This leads to the desired result. □

What remains is to determine w and to find a means of computing x_h. Note that upon substituting (5) into $A'x = b'$ one obtains $A'(A')^* w = b'$, which upon using (1) becomes $\mathcal{I}_m^S \; w = b'$. This can be used to re-express equation (5) as

$$\begin{aligned} x &= (A')^* \; w + x_h \\ &= (A')^* \mathcal{I}_m^S \; w + x_h \\ &= (A')^* \; b' + x_h \end{aligned}$$

which is the same as in equation (4).

Discussion

Of note is that $x_p = Gb = (A')^* b'$ is a minimum norm solution. This follows since it is formed from row space vectors, and not null space vectors. Its minimum norm nature is also evident since the generalized inverse is, in general, of type $\{124\}$ (see Appendix).

There are two variations to this approach, which are mainly of interest for numerical applications. When G and M are not explicitly computed, the solution is:

$$\begin{aligned} x_p &= (A')^* b' \\ x_h &= Py \\ P &= 1_n - (A')^* A' \end{aligned}$$

where again y is an arbitrary vector in $\mathcal{F}^n$. On the other hand, if G is explicitly computed, the solution may be written as:

$$\begin{aligned} x_p &= Gb \\ x_h &= Py \\ G &= (A')^* M \\ P &= 1_n - GA \end{aligned}$$

where y is the same. Depending on the application, it may or may not be advantageous to explicitly compute G.

4 Column-Based In Situ Approach

In this approach, it's first recognized that $Ax = b$ will have a solution only when b lies in the column space of A. In particular, a least squares solution is sought, in which only the portion of b that lies in that column space is allowed; this is defined as b_c. Hence, the equation to be solved is $Ax = b_c$, and the strategy now is to solve it by attempting to orthonormalize the columns of A by right-multiplying it by M_s. Thus, matrix products $M_s M_s^{-1}$ are repeatedly inserted between A and x in $Ax = b_c$, to produce

$$A(M_1 M_2 \cdots)(\cdots M_2^{-1} M_1^{-1})x = b_c$$

or

$$A' M^{-1} x = b_c$$

where $M = (M_1 M_2 \cdots)$ and $A' = AM$.

As explained in the section Definitions, the matrices M_s will effect an orthonormalization procedure on the columns of A. Observe that A' can be used in a projection operator to compute $b_c = A'(A')^* b$, and the equation to be solved can now be written

$$\begin{aligned} A' M^{-1} x &= A'(A')^* b \\ A[x - M(A')^* b] &= 0 \end{aligned}$$

The piece in the square brackets must lie in $\mathcal{N}[A]$; it is named x'_h. Note that the solution x and vector x'_h are now entwined in the equation $x - M(A')^* b = x'_h$ (and one may be used to define to define the other, depending on which is taken as given). Finally, observe that x'_h may be defined to be any null space vector, and that it will simply shift the solution x. A solution is now proved which assumes an arbitrary null space vector x_h.

Theorem: A solution to $Ax = b_c$ is

$$x = M(A')^* b + x_h \tag{6}$$

where the variable definitions from this section are used.

Proof. The proposed solution is verified by substitution into Ax

$$\begin{aligned} Ax &= AM(A')^*b + Ax_h \\ &= A'(A')^*b \\ &= b_c \end{aligned}$$

□

Discussion

In this column-based approach, the set S is defined to be the columns of A' that are nonzero, so that $\mathcal{I}_n^S$ is the index matrix associated with A'. It is easy to verify:

$$A' = A'\mathcal{I}_n^S \tag{7}$$

$$(A')^* = \mathcal{I}_n^S(A')^* \tag{8}$$

These identities are useful in determining that the generalized inverse $G = M(A')^*$ is of type $\{123\}$ (see Appendix). In summary, the (least squares) solution found for this section can rewritten as

$$\begin{aligned} x_p &= Gb \\ x_h &= Py \\ P &= 1_n - GA \end{aligned}$$

where, once again, y is an arbitrary vector in F^n.

5 Online Capabilities

It is normally assumed that all data (i.e., A, b) is present before solving $Ax = b$ may commence, using any of the usual methods. For the algorithms introduced here, it will be seen this is *not* required. Indeed, the linear solver may begin its work when the data arrives in a certain regular manner; this is referred to as an *online* algorithm. This capability may find use in the case where the dimensions of A are large, and it is also a time-consuming process to compute its entries.

Row-Based

Assume that the rows of $[A|b]$ become available one at a time to the solver, in the order 1 to m. The orthonormalization steps described earlier may be done incrementally as these rows become available. When the GS procedure is used to effect the orthonormalization, the 1-st through i-th rows of A, M and b will no longer change following the i-th step of the algorithm. It follows that these rows can be used in a computation of x_p. Taking advantage of the column-row expansion [5], x_p may be written as

$$x_p = (A')^* b' = \sum_{i=1}^{m} x_p^{(i)}$$

where

$$x_p^{(i)} = \mathrm{Col}_i[(A')^*]\ b'_i$$

and Col_i signifies the i-th column. Thus following the i-th step in the algorithm, $x_p^{(i)}$ may be computed. Also note that the updates to x_p are mutually orthogonal, i.e., $< x_p^{(i)}, x_p^{(k)} >= 0$ for $i \neq k$. This means that the estimation of $\|x\|$ is monotonically non-decreasing as more rows are included.

One can take a similar approach when the second variation is used. For example, the generalized inverse may be computed in an online mode as:

$$G = (A')^* M = \sum_{i=1}^{m} \mathrm{Col}_i[(A')^*]\ \mathrm{Row}_i[M]$$

Note that the summand may be computed following the i-th step of the GS procedure. Following the computation of G, the particular solution is easily computed from it.

Column-Based

Here it will be assumed that all of b is available at the outset, and that the columns of A become available one at a time to the solver, in the order 1 to n. When the GS procedure is used to effect the orthonormalization, the 1-st through j-th columns of A and M will no longer change following the j-th step of the algorithm. Since these columns in M and A are done changing at these points, they become available to be used in a computation of x_p.

Again, a column-row expansion [5] may be used to rewrite x_p, this time as

$$x_p = \sum_{j=1}^{n} x_p^{(j)}$$
$$x_p^{(j)} = \mathrm{Col}_j[M]\ \mathrm{Row}_j[(A')^*]\ b$$

Since the j-th row of $(A')^*$ and j-th column of M are available at the j-th step, $x_p^{(j)}$ can be computed at the j-th step. Also observe that this same online technique can be trivially extended to compute G, P or even x_h in an online manner.

Complexity of Online Approaches

At play are two time scales: the time needed (τ) for a new portion of the data to be acquired, and the time needed (Δ) for the linear solver to process the new portion of data. These two timescales determine the additional complexity due to the solver, which by definition will not include the time spent waiting for the data to arrive. If $\tau > \Delta_{max}$, the solver will generally be waiting on data acquisition. If $\tau < \Delta_{min}$, the complexity will be the same as if all data is available at the outset.

For the row-based approach, the time used by the solver on the i-th row is $\Delta(i) \sim O(4ni - n)$, which varies from $\Delta_{min} = \Delta(1)$ to $\Delta_{max} = \Delta(m)$. Also, no matter how large the size of τ, the solver will always add complexity due to the computations done on the last row ($i = m$), as well as the final matrix-vector multiply (i.e., forming $(A)^*b$). Combined, these two contribute a complexity of $O(4mn)$, which acts as the minimum amount of added complexity for the row-based solver. The maximum complexity occurs when $\tau < \Delta_{min}$, and in this case is $O(2m^2n)$. Thus the complexity added by the solver is:

$$O(4mn) < \text{complexity} < O(2m^2n)$$

Similarly, in the column-based approach, the time used by the solver on the j-th column is $\Delta(j) \sim O(4mj - m)$, which varies from $\Delta_{min} = \Delta(1)$ to $\Delta_{max} = \Delta(n)$. By a similar set of arguments, the complexity added by the solver in this case is

$$O(4mn) < \text{complexity} < O(2mn^2)$$

Thus the acquisition time τ acts as a crossover parameter for the complexity of the algorithm.

6 Final Remarks

The approaches discussed herein, which came about while working on the inverse kinematics problem in robotics, uses a slightly different view on the QR factorization method. Instead, the factorization is done when A is in the context of the equation $Ax = b$. Borrowing a term used in metallurgy, it is described as an "in situ" approach.

One of the new features the algorithms present is that it can be computed in an online manner. This is expected to be useful in applications where the data is large, or when it simply takes a long time to create and/or acquire the data. What is also immediately noteworthy about the algorithm is that it doesn't use an elimination or partitioning strategy; in particular, there was no solving of a triangular system of equations.

The two in situ approaches have a symmetry between them. The (second variation of the) row-based in situ method involves a transformation of $[A|1]$ to produce a particular solution.

$$[A|1] \longrightarrow [A'|M] \longrightarrow x_p = (A')^* M\, b$$

Likewise, the column-based in situ method may be expressed as:

$$\begin{bmatrix} A \\ 1 \end{bmatrix} \longrightarrow \begin{bmatrix} A' \\ M \end{bmatrix} \longrightarrow x_p = M(A')^*\, b$$

The main reason for mentioning it is aesthetics: the two solutions comprise a more symmetrical viewpoint on solving $Ax = b$.

Although the new algorithms were cast to solve $Ax = b$, they can easily solve its matrix generalization [13]: $AX = B$, where X is n-by-p and B is m-by-p (and $p \geq 1$). In that case the orthonormalization steps proceed as before, and the particular and homogeneous parts of the solution $X = X_p + X_h$ are

$$X_p = GB$$
$$X_h = PY$$

where G and P are the same as before, and $Y \in \mathcal{F}^{n \times p}$. This also admits an online formulation.

Acknowledgements

The author dedicates this paper to Robert E. Zimmer for his kind support and encouragement. In addition, the author thanks Prof. Daniel Grayson for a discussion of the manuscript.

APPENDIX - Penrose Conditions

It is convenient to classify a generalized inverse (G) according to which Penrose conditions it satisfies ([9, 10, 4, 2, 1]). In particular, different properties of the solution Gb follow if certain sets of the following Penrose conditions are satisfied.

$$
\begin{aligned}
&(1) AGA = A \\
&(2) GAG = G \\
&(3) AG = (AG)^* \\
&(4) GA = (GA)^*
\end{aligned}
$$

Row-Based In Situ

Here the index matrix will be $\mathcal{I}_m^S$, where the set S is the set of non-zero row vectors in the matrix A', as defined in the section "Row-based In Situ Approach". The identities involving $\mathcal{I}_m^S$ are used herein.

The first Penrose condition is seen to always be true for our solution.

$$
\begin{aligned}
AGA &= A(A')^*MA = M^{-1}MA(A')^*A' \\
&= M^{-1}A'(A')^*A' = M^{-1}\mathcal{I}_m^S A' \\
&= M^{-1}A' = A
\end{aligned}
$$

Likewise, the second condition is always true.

$$
\begin{aligned}
GAG &= [(A')^*M]A[(A')^*M] \\
&= (A')^*A'(A')^*M = (A')^*\mathcal{I}_m^S M \\
&= (A')^*M = G
\end{aligned}
$$

The third Penrose condition is seen to be problematic

$$\begin{aligned} AG &= A(A')^*M = M^{-1}MA(A')^*M \\ &= M^{-1}A'(A')^*M = M^{-1}\mathcal{I}_m^S M \end{aligned}$$

If A is of full row rank, then $\mathcal{I}_m^S$ equals 1_m, AG becomes the identity matrix 1_m, and the condition is manifestly satisfied. However, if A is *not* of full row rank, then this condition is not in general satisfied. Finally, the fourth Penrose condition is seen to be true:

$$\begin{aligned} GA &= (A')^*MA = (A')^*A' \\ &= ((A')^*A')^* = ((A')^*MA)^* \\ &= (GA)^* \end{aligned}$$

In summary, the generalized inverse obtained by this algorithm is at least a $\{124\}$-inverse. If it's additionally true that A is of full row rank, then it becomes a unique $\{1234\}$-inverse.

Column-Based In Situ

Here the index matrix will be $\mathcal{I}_n^S$, where the set S is the set of non-zero column vectors in the matrix A', as defined in the section "Column-based In Situ Approach". The identities involving $\mathcal{I}_n^S$ are used herein.

The first Penrose condition is seen to always be true.

$$AGA = AM(A')^*A = A'(A')^*A = A$$

To understand this, recall that $A'(A')^*$ projects away vectors not in $\mathcal{R}(A)$; hence it leaves A unchanged. Likewise, the second condition is always true.

$$\begin{aligned} GAG &= [M(A')^*]A[M(A')^*] \\ &= M(A')^*A'(A')^* = M\mathcal{I}_n^S(A')^* \\ &= M(A')^* = G \end{aligned}$$

The third Penrose condition is also true

$$\begin{aligned} AG &= AM(A')^* = A'(A')^* \\ &= [A'(A')^*]^* = [AM(A')^*]^* \\ &= (AG)^* \end{aligned}$$

The fourth Penrose condition $GA = (GA)^*$ is seen to be problematic

$$\begin{aligned} GA &= M(A')^*A = M(A')^*AMM^{-1} \\ &= M(A')^*A'M^{-1} = M\mathcal{I}_n^S M^{-1} \end{aligned}$$

Observe that if A is of full column rank, then $\mathcal{I}_n^S$ equals the identity matrix 1_n, and GA becomes 1_n; this trivially allows the identity to be satisfied. However, if A is *not* of full column rank, then this condition is not true in general.

In summary, the generalized inverse obtained by this algorithm is at least a $\{123\}$-inverse. If it's additionally true that A has full column rank, then it becomes the unique $\{1234\}$-inverse. Recall that generalized inverses which are at least of type $\{13\}$ (which is the case here) provide solutions Gb with a least-squares property.

References

[1] R. B. Bapat, *Linear Algebra and Linear Models, 2nd ed.*, Springer, 2000.

[2] A. Ben-Israel and T. N. E. Greville, *Generalized Inverse, Theory and Applications, 2nd ed.*, Springer-Verlag, New York, 2003.

[3] R. M. Bowen and C.-C. Wang, *Introduction to Vectors and Tensors, 2nd. ed.*, Dover, 2008.

[4] S. L. Campbell and C. D. Meyer, *Generalized Inverses of Linear Transformations*, SIAM, 2008.

[5] D. Carlson, C. R. Johnson, D. C. Lay, and A. D. Porter, eds., *Modern views of matrix multiplication*, pp. 5–8, in *Linear Algebra Gems*, The Mathematical Association of America, 2002.

[6] L. Giraud, J. Langou, and M. Rozloznik, CERFACS Technical Report No. TR/PA/02/33. http://www.cerfacs.fr/algor/reports/2002/TR_PA_02_33.pdf

[7] G. H. Golub and C. F. Van Loan, *Matrix Computations, 2nd ed.*, The Johns Hopkins University Press, Baltimore, 1989.

[8] B. Noble, *Methods for computing the Moore-Penrose generalized inverse, and related matters*, in *Generalized Inverses and Applications*, M.Z. Nashed (ed.), Academic Press, New York, 1976, pp. 245–301.

[9] R. Penrose, *A generalized inverse for matrices*, Proc. Cambridge Philos. Soc., **51** (1955), 406–413.

[10] R. Penrose, *On best approximate solutions of linear matrix equations*, Proc. Cambridge Philos. Soc., **52** (1956), 17–19.

[11] G. W. Stewart, *The decompositional approach to matrix computation*, Computing in Science & Engineering, Jan/Feb 2000, pp. 50–59.

[12] D. S. Watkins, *Fundamentals of Matrix Computations*, John Wiley & Sons, New York, 1991.

[13] M. F. Zimmer, private notes.

[14] As an implementation detail, note that b may be changed in place to become b', or that the changes may be stored in M and b' can be later computed as $b' = Mb$.

[15] Note that the GS procedure can be implemented as either Classical (CGS) or Modified (MGS). An analysis of their stability properties while using reorthogonalization is done in [6].

[16] Note that GS procedures can be formulated to be forward or backward looking, with respect to the list of vectors that a new vector is orthonormalized against. For these online approaches, the backward looking approaches are assumed.